

This report contains a detailed description of the inner workings of the Time-Sharing Director. The report is provided as 'support documentation' in the sense described in section 2 of the KDF9 Service Routine Library Manual.

The report should be read in conjunction with the KDF9 (which can be obtained from the current MINISTER - see Service Routine Library Manual) and the corresponding report on the Non-Time-Sharing Director. References are made to section A of the latter - by the letter 'N' followed by the appropriate section number.

The Director described is the Standard Time-Sharing Director and the report takes no account of optional extra facilities.

The Company reserves the right to change the information contained in this report without notice.

LIBRARY SERVICES,
SYSTEMS PROGRAMMING DEPARTMENT.
1ST MAY, 1965.

CONTENTS

	Page
1 Hardware Features	1
1.1 General	1
1.2 Priorities	2
1.3 The Program Holdup (PHU) store	2
1.4 Program Ready RFI	4
1.5 Switching of Nests, etc.	5
1.6 Summary of Director-only instructions	7
2 Storage and input of programs	8
2.1 A and B levels	8
2.2 Storage assignment	9
2.3 Housekeeping in the Director	10
2.4 Program movement and Priority interchange	12
2.5 Program Input and Termination	13
3 Interruption Processing	16
3.1 Path through the Director	16
3.2 The Short Path	17
3.3 The Long Path	18
3.4 Timing	20
3.5 V-stores of P0	20
4 Hold-ups and Subprograms	21
4.1 Type of hold-up. Action following EDT	21
4.2 FLEX. The action of P4.	32
5 OUTs (P3, etc.). Program Failure (P2)	36
5.1 P2 (Program Failure)	37
5.2 OUT Routines	38
5.3 P23 and OUT 8	42
6 Subroutine Index	45

- 1 -

1. Hardware Features

1.1 General

Before describing the special hardware features of the Time-Sharing KDF9, it details a brief account of their objectives is necessary. The standard (non-Time-Sharing) KDF9 is already provided with interruption facilities which enable the operating program to be interrupted automatically, and thus to pass control to the Director, when the program is held up by the action of a peripheral transfer - either because it is linked to a specific peripheral unit involved because it requires access to the area of core store which is locked-out for the duration of the transfer. However, there is no hardware which will record automatically which peripheral unit or which part of the core store, is causing the hold-up. Nor is there any indication given at the end of a particular transfer that was the cause of a hold-up.

The object of "time-sharing" in KDF9 is to be able to have more than one program (besides the Director) stored in the machine at a time, and to run them on a "priority" basis, so that the highest priority program runs until it is held up (or else), when control will be passed to the Director, who then moves it in turn, runs until it is held up or ends, or until the conditions which caused any higher priority program to be held up are removed; and so on. Thus, whenever a program is interrupted, one expects the Director to return control to one of the other programs, rather than waiting until the interrupted program can resume.

The existence of the KDF9 "Base Address" (N1.1) and the store and peripheral protection facilities provide by NOL (N2.1.3) and CPDAR (N2.1.4) make the storage (and movement) of several programs a safe and comparatively simple process. However, to implement efficiently the priority system described above, four extra hardware facilities are required, which makes the difference between a "Time-sharing" and a "non-time-sharing" KDF9:-

- (a) a method of associating a priority with each program;
- (b) a "Program Hold-up" store to record in a dynamic manner which programs are held up at any time, and why;
- (c) an RFI which will cause a program, to be interrupted on the removal of a holdup condition which was causing any program of higher priority to be held up;
- (d) a method of speedily interchanging the contents of the Nest, SJNS and Q-stores, for any two programs, so that the transfer of control from one program to another is as fast as possible.

The hardware implementation of these facilities, and the instructions which use them, are described below.

- 2 -

1.2 Priorities

whenever any program is running it has associated with it a "current priority level" (CPL) in the range 0-3. This number is kept in a 2-bit register which is set by the instruction = K1 (N 1.2, N 2.1.3) which besides transferring D38-47 of N1 to BA and D24-33 to N0, sets CPL equal to the value given in D34, 35 of N1.

Priority 0 is the highest priority, priority 3 the lowest.

It would be possible for several co-existent programs to have the same priority but this contradicts the principles underlying the design of the PHU store (see 103 below) and the multiple Nesting stores, so practice there are up to 4 programs, each with its own priority number.

It is customary to refer to "priority n" as meaning (in the right context): "the program which runs at priority level n". It is possible, subject to certain practical restrictions to alter the priority of a program.

The Director itself customarily runs with priority 0 but, when the Director is about to execute a peripheral transfer on behalf of a program (i.e. involving that programs core area) CPL must be set to the value of that programs priority.

The actual significance of CPL in two-fold: firstly, any peripheral transfer has associated with it for its entire duration the value of CPL when it was called (this affects what happens when the transfer ends); and secondly, it determines whether or not the PR RFI (1.4) occurs when a peripheral hold-up is removed.

1.3 The Program Hold-up (PHU) store

This is a set of four 12-bit (D0-11) registers called PHU0, PHU1, PHU2 and PHU3, which are used to record, dynamically, peripheral holdups incurred by priorities 0,1,2 and 3 respectively.

When priority n is interrupted because of a peripheral hold up (LOV), PHU n automatically records the hold-up as follows: D11 is set to 1, and D10 is set to 1 or 0 according to whether the hold-up is due to a reference to a busy peripheral buffer, or to a locked-out store. In the former case, the busy buffer number is stored in D6-9 (D0-5 are cleared), in the latter case the most significant 10 bits of the 15-bit absolute address of the locked-out location are stored in D0-9. This is known as the "group address".

The PHU store is not altered if LOV occurs in Director mode, i.e. if the Director refers to a busy buffer or locked-out area.

- 3 -

PHU n is also altered if priority n executes the instruction CLOQ, referring to a busy buffer. Again D11 and D10 are set to 1, the buffer number is stored in D6-9, and D1-5 are cleared; but D0 is set to 1.

Note that D11=1 in PHUn implies that priority n is held up by a peripheral transfer.

The end of a peripheral transfer which was called by priority n (i.e. called at any time, in Director or program Mode, when CPL had the value n) will cause PHU n to be cleared if and only if,

- (a) D0, 10 and 11 of PHU n are all non-zero
- or (b) D10 and 11 are non-zero and the buffer whose number is given by D6-9 is not busy
- or (c) D11 is non-zero, D10 is zero, and D0-9 give the absolute group address of a block of 32 words which is not locked-out.

This represents the removal of hold-ups due to (a) "Interrupt if Busy" - the holdup being removed at the end of any transfer called by priority n; (b) reference to a busy buffer; (c) reference to a locked-out part of the store.

In simple terms, PHU n is set automatically when priority n is held up, and will be cleared automatically at the end of the appropriate peripheral transfer, provided that that transfer was called by, or on behalf of the same priority, i.e. when CPL was equal to n.

It is possible to ensure that no program ever refers to a buffer which has been made busy by another program. An obvious example is the "shared" magnetic tape buffer, which is made busy when any one of the units connected to it - which might be allocated to different priorities - is carrying out a transfer. Therefore, in order to avoid a situation where a program being transferred is held up for ever, the hardware is so arranged that at the end of any peripheral transfer called by priority n, not only is PHU n examined and, if need be, cleared, but also the remaining 3 PHU registers are examined, and if any of them are showing "busy" hold-ups due to a buffer which is no longer busy, the EDT RFI is set, and will of course be set anyway if the transfer was called in Director mode.

- 4 -

Note that these other PHU registers are not cleared, in this case. To do so would be dangerous: it is preferable that the Director should be more positively aware of the situation, in order that it can do whatever is necessary to run the system.

The priority rule is, in this case, to allow one of the lower priority programs to use the shared buffer, for instance. In a case where two programs are sharing a buffer, and the higher priority of the two makes frequent use of it while the lower priority only uses it occasionally, a judicial reversal of priorities, to allow the lower user to have priority access to the buffer in preference to the higher, can lead to more efficient operation.

In order to have the necessary control over the PHU store, the Director must be able to examine its contents, and to clear the individual parts of it.

The instruction K5 causes the least significant 6 bits of each part of the PHU store to be brought into N1, nesting down one place. Thus,

D6-11 of PHU0 occupy D0-5 of N1
 " " " PHU1 " D6-11 " "
 " " " PHU2 " D12-17 " "
 " " " PHU3 " D18-23 " "

Thus the Director can tell which priorities are held up at any time, which are held up by busy buffers, and in such cases, by which buffers.

The instruction CLOQ, as well as clearing lock-outs from the specified area, also causes PHU0 to be cleared, where n is the value of CPL when the instruction is obeyed.

Note that the Director can determine when a program is held up by a lock out, but not the address involved (because it only sees 4 bits of this address).

The effect of all this is that Director must always look at the PHU store to determine which priorities can be resumed, i.e. are not held up; and also, whenever EDT occurs, Director must see whether any parts of the PHU have to be cleared, and, if so, whether the normal priority system should be over-ruled.

1.4 Program Ready RFI

This is the end of a peripheral transfer called by priority n, the RFI is expected to look at D11 of each part of the PHU store to see which priority can now resume.

- 5 -

In practice this means that if the hardware detects the end of a peripheral hold-up on a program which is of higher priority than the one currently operating, an interruption occurs.

This RFI corresponds to D22 in the copy of RFI which the instruction K4 (N2.1.5) transfers to N1. When this RFI is detected the Director is expected to look at D11 of each part of the PHU store to see which priority can now resume.

LOV corresponds to D28 in N1, after K4. In this case a change to a lower priority level can be expected.

The normal setting of PR is inhibited when the transfer whose ending led to the clearing of the PHU register is also found to have been the reason for "buffer busy" hold-ups at other priorities: in such cases EDT is set anyway, as described in 1.3 above.

PR is also set by the instruction "Interrupt if Busy" (INTQ) when the relevant buffer is found to be busy, in order to cause the required interruption. This avoids the need to have a separate RFI for this case. The same action is required of Director as in the normal PR or LOV interruptions.

1.5 Switching of Nests, etc.

To provide the necessary rapid takeover of the contents of Q-stores, Nests and SJNS, four Sets of each are provided on the Time-Sharing KDF9, together with instructions for switching from one to another.

The Sets are numbered 0-3 but this numbering is completely independent of the priority numbering system. Each Set comprises 15 Q-stores (1-15), a Nesting-store stack of sixteen 48-bit cells, and an SJNS stack of sixteen 16-bit cells. The two stack counters are not quadruplicated, nor are the overflow and Test registers.

In order to switch from one Set to another, it is necessary to obey the instruction = K3 with the following parameter in N1:

D0,1 = New Set number
 D2-6 = Nest stack counter
 D7-11 = SJNS stack counter

The effect of all this is that the Director constantly inspects the store assignments to ensure that programs are packed tightly into the available space.

When adding corrections to the "stack counters" in N1 (as fetched by K3), it is important to ensure that no spill occurs from one counter to the other, or into the Set number position. This could only happen if NOV had occurred.

The "Set parameter" used in the example is the record of the Set number, and the stack counter initially used by it. It consists of two bytes of 3 and 1 respectively, in the Nest and SJNS stack counters over the values these took just after the program they pertain to was interrupted.

1.6 Summary of Director-only instructions

This includes all the instructions of this type used by the Time-Sharing Director. Some of these have greater power than they do in a non-TimeSharing KDF9.

1.6.1 EXITD (N 1.1) As well as clearing the lockout specified by INTQ, this instruction also clears PHU0, where n is the current value of CPL

1.6.2 CLOQ (N 2.1.2) Clear transfer

1.6.3 CTQD (N 2.1.2) Clear transfer or off. There is a "buzzer" attached to PHU0 which can be switched on or off by D6-9. If the value of D6-9 is non-zero, then the buzzer will be turned on. If the value of D6-9 is zero, then the buzzer will be turned off.

1.6.4 =K0. Switch buffer or off. There is a "buzzer" attached to PHU0 which can be switched on or off by D6-9. If the value of D6-9 is non-zero, then the buzzer will be turned on. If the value of D6-9 is zero, then the buzzer will be turned off.

1.6.5 =K1. (N2.1.3) This instruction transfers D24-33 of N1 to BA, and D34-47 to N0. No nesting

1.6.6 =K2. (N2.1.4) Set CPDAR. No nesting

1.6.7 =K3. Switch one Set of Q-stores, nests, etc., using D0-11. Nest stack counter

(D0,1 = New Set number
 D2-6 = Nest stack counter
 D7-11 = SJNS stack counter)

1.6.8 K4. (N2.1.5) Fetch PHU0 and clock. Two significant bytes of the stack counters before this nesting operation has been recorded.

So to get a true record of the state of the counters after obeying K7, it is necessary to add D6 to N1 -- the equivalent of adding 1 to the Nest stack counter.

1.6.9 K5. (N2.1.6) Set CPDAR (D28), Nest down 1 place.

K5. D6-11, D12-17, D18-23, respectively, of N1. Nests down 1 place.

1.6.10 K7. (N2.1.7) Set CPDAR (D28), Nest down 1 place. The nest counter value to N1, in format as for K3. Note that the value of the Nest stack counter shown in D2-6 is 1 less than the value it actually had after obeying this instruction, because the nesting down 1 place.

- 6 -

Operations (e) and (f), above, would in practice probably refer to locations which depend on priority numbers, so the addresses used are likely to be modified by the contents of Q-stores. Care must be taken to ensure that those Q-stores have the correct values (their contents before interruption) restored to them before =K3 is obeyed.

1.7 Storage Assignment

The part of the core store which a program occupies - always a continuous area, and a multiple of 32 words beginning at an absolute address which is also a multiple of 32 - is said to be "assigned" to the program. The total area available for assignment to programs extends from the very top (highest address) of the store, down to the lowest word in the store which is not part of the Director, and whose address is a multiple of 32.

"A" programs are stored at the top of this area, and B programs at the bottom, next to Director. The Director constantly inspects the store assignments to ensure that programs are packed tightly into the available space.

It is necessary for the Director to physically move programs in the store to take over space vacated by other programs, and to assign new areas to them.

It is also necessary for the Director to physically move programs in the store to take over space vacated by other programs, and to assign new areas to them.

The Director must always look at the store assignments to ensure that programs are packed tightly into the available space.

The Director must always look at the store assignments to ensure that programs are packed tightly into the available space.

The Director must always look at the store assignments to ensure that programs are packed tightly into the available space.

The Director must always look at the store assignments to ensure that programs are packed tightly into the available space.

The Director must always look at the store assignments to ensure that programs are packed tightly into the available space.

The Director must always look at the store assignments to ensure that programs are packed tightly into the available space.

The Director must always look at the store assignments to ensure that programs are packed tightly into the available space.

The Director must always look at the store assignments to ensure that programs are packed tightly into the available space.

The Director must always look at the store assignments to ensure that programs are packed tightly into the available space.</p

This required amount of store, from the upper limit of the Hole downwards, is assigned before starting to read in the A program (unlike the procedure described above for B programs, which starts with only 32 words assigned). Once the "B" block has been read in, this area may be contracted (from the top downwards) if the actual store requirement is less than the amount assigned; the vacant space thus created will be transferred to the "Hole in the middle" by the normal program-moving procedure.

Note that core storage ceases to be "vacant" as soon as it has been assigned for program input, and is therefore liable to be "slid" up or down the store while the program to occupy it is still being input. However, an assigned area can only be moved when there are no peripheral lock-outs set on it.

The case described above, when the store area assigned to a B program "expands", between reading the B and C-blocks, from 32 words to the requisite number, is the only one which ever involves increasing the amount of store assigned to a program; it is the only one where there is a "vacant" area, i.e. the "hole in the middle" to expand into. The area assigned to a program may be decreased on three occasions: firstly, in the case already described, where an A program's store requirement is less than the amount pre-assigned for reading it in; secondly, when a program obeys OUT 1; and thirdly, when a program obeys OUT 2. In the last two cases, the new program/section and its predecessor have the same ED, and the contents of the program/section in both old and new areas remain unaltered (provided of course, that in the case of OUT 1, they have not been overwritten by the latter). But the new program/section must not demand more store space than the old; if it does, failure occurs. If the new program/section has a blank store requirement, it is given exactly the same assignment as its predecessor.

The new program/section has the same priority, and the same A/B characteristic, as the old one, and, for an A program, is pre-allocated the same peripheral units.

Only one A and one B program can be in the process of input at any time - this includes the input of a new program/section after OUT 1, which may therefore have to wait if program input is already going on at another priority.

2.3 Housekeeping by the Director

Whenever any program is loaded it is assigned one of the four Sets of stacks and Q-strokes. Although a program may change its priority whilst it is in the machine, it always keeps the same Set of stacks, etc., and it is convenient to identify each program by the name of its "Set". To avoid confusion with priority numbers, the "Sets", and programs, are called P, Q, R, and S. The bottom two bits of the octal representation of these characters - 60, 61, 62, 63 - give the actual number of the "Set".

Inevitably, a great many of the "housekeeping" parameters used by the Director are in groups of 4 consecutive words (quartets), one word for each program. Within each quartet the words are arranged in priority number order, so that the first word pertains to priority 0, the second to

priority 1, and so on. This arrangement is chosen (in preference to Program letter order) because in practice the priority number occurs more naturally and hence is a more convenient parameter for indirect addressing. In almost every case in which the priority number is used for indirect addressing (or for any other purpose) is is treated as M5, and so the conventions adopted in this respect of referring to e.g. V29P0M5 as meaning the 4 words stored in V29P104, V30P104, V31P104 and V32P104, which have similar significance in relation to priorities 0, 1, 2 and 3, respectively.

The housekeeping parameters which are particularly relevant to the storage of programs in the machine are the single word VOP104, and the quartet V9P0M5.

The most significant 24 bits of VOP104 are assigned in groups of 6 to the four priorities so that D0-5, D6-D11, D12-17 and D18-23 correspond respectively to priorities 0, 1, 2 and 3. The bits for priority 0 have the following significance:

```
D5 = 1 if priority 0 cannot be entered for any reason (apart from a peripheral hold-up)
D4 = 1 if priority 0 is a pre-assigned "A" level
D3 = 1 if priority 0 is occupied by an "A" program
D2 = 1 if priority 0 is occupied by a "B" program. Obviously D2 and D3 cannot both be non-zero; and if D3 = 1, D4
```

must be non-zero.

The corresponding bits of the other groups of six have the same significance in relation to priorities 1, 2 and 3.

The "hold-up" bits, D5, 11, 17, 23, allow for every possible "software" reason for not entering a priority, and could indicate, for instance, that the current program is awaiting a certain priority level, or merely that a priority has held up pending the allocation of a peripheral unit. Note that after obeying the instruction K5 (1.3, 1.6.9), bits D5, D11, D17, D23 of N1 indicate peripheral hold-ups; and therefore the sequence of instructions

```
VOP104; K5; OR;
```

will leave a complete indication in those bits of N1, of which priorities can be entered and which cannot.

The quartet V9P0M5 contain the Base Address, Number of Locations and Priority Number for each priority. Just before any program is entered the contents of the appropriate word is transferred, by executing =K1 (1.6.5) to the BA/NOI/EPI registers.

The quartet V9P0M5 therefore indicate the areas of store occupied by each priority. This includes storage which may be in the process of being "expanded" or "contracted" during program input - it may, for instance, amount to only 32 words, during the input of the B-block of a B program. The A portion of any word of this quartet is zeroed if, and only if, there is no vestige of a program at that particular priority level. It is only when BA = 0 that D2 and D3 of the corresponding 6-bit group in VOP104

are both zero: as soon as a priority is assigned any storage at all, it acquires the status of either an A or a B program, and keeps the status just as long as it has a foothold in the store; i.e. D38-47 (BA) of V9P0M5, and D2-3 of the appropriate group in VOP104 are set and cleared together.

2.4 Program movement and priority interchange

One of the properties of KDF 9 Time-sharing is that programs may have their priorities interchanged, and may be moved about in the store.

The priority of a program may be changed either as the result of a TINT V, or, following the priority upgrading which may occur at the end of an A program. A program may be moved in the store as a result of the ending of a program leading to the expansion of the "Hole in the Middle".

As explained in 4.3 above, it is dangerous to change the priority of a program whose part of its area is locked out - because it may be a peripheral hold-up getting "stuck" in the wrong PHU register. Similarly, a program may not be moved in the store while any part of its area is locked out - a lock-out implies that a transfer is going on in that area and so the contents of the area are liable to change. Therefore if for any reason for moving a program, or changing its priority, occurs while any part of the program's area is locked out, the details of the re-order, change are recorded in the appropriate member of the quartet V9P0M5, and at the same time the program is temporarily prevented from resuming, using the appropriate bit in VOP104.

The impending change is recorded in V9P0M5 as follows:

```
D0-2 = quantity to be added to priority number
D14-23 = new value of BA.
```

The program is only allowed to resume when D0-23 are all zero.

Not only is the program held up when a priority change or store move is impending, but also the Director must be prevented from starting any part of the program which would interfere with the transfer and prolong the duration of the lock-out on the area. Therefore before any peripheral transfer executed by Director, a check is performed that the transfer area does not lie within a region belonging to a program which is going to be moved or has its priority altered. Subroutine P103 performs this check, requiring as data the (absolute) addresses of the transfer area, the subroutine also performs the important function - that the transfer is to be all or part of a program. P103 also takes care of CPL to that pertaining to the program whose area is involved. This meets a requirement mentioned in 1.3; again, it prevents the probability of a peripheral holdup getting "stuck".

When the priorities of two programs are interchanged the relative positions of the relevant numbers of all the parameter quartets have to be changed accordingly. All the quartets which have to be re-ordered

in this way are stored in the V-stores of P104, from V1 onwards. There are two other quartets which require a slightly more complicated re-arrangement: one is the quartet V9P0M5, in which the values of the

The manual reads V9P0M5 but this has to be type

priority number (D4-5) and the priority displacement (D0-2) have to be modified as well; the other is the quartet V29P104M5, which comprises the four "TAB" words, each containing the characters CR-LF, Case Normal, then sufficient TABs to ensure that any message typed out for the relevant program starts in the right column, the priority number and a spaces here the priority number (D40-1) has to be modified. In addition, the appropriate subroutines have to be altered. Any priority number recorded in a specific location (e.g. the priority level at which A or B program input is taking place) and the priority number of the last program interrupted) must be altered if necessary. The priority-dependent subprogram hold-up table (4.1.5) and subprogram numbers in the various peripheral transfer queues (4.1.3) also have to be modified appropriately. All these functions are performed by the subroutine P105, which carries out all priority changes and store moves, and also continually inspects the store and the state of the A programs to see what new moves and priority changes need to be instigated.

When a program's area is moved in the store (also by P105), the appropriate member of the quartet V9P0M5 has to be modified (D14-23 and D38-47). Also any "absolute" addresses, in the subprogram tables and the various peripheral transfer queues (see 4.1.05 and 4.1.3), which are affected, have to be altered. This is all looked after by P105, which is entered after every EDT.

The parameters contained in the various quartets of P104 are defined in the User-code listing. One quartet, V25P104M5, is of particular interest here. The offset of the subroutines concerned is to any priority in VOP104 set to 4, and only if, the corresponding member of the quartet is non-zero. The various bits in each word of this quartet when non-zero, each indicate a different reason for a holdup, as follows:-

```
D43 = 1 implies a store move and/or priority change impends
D44 = " a hold-up due to the "OUT 8" subprogram
D45 = " the program is suspended by TINT 8
D46 = " the program is absent, or being input or terminated
D47 = " the main subprogram is active
```

The significance of D44 and 47 is explained in the sections dealing with subprograms (4.1.5, 4.1.6, 5.3). The termination hold-up, D46, is described in 2.5, below.

The subroutines P113 and P114 are used respectively, to set and clear bits in V25P104M5, and at the same time set or clear the hold-up bits in VOP104, as appropriate.

2.5 Program Input and termination

While a program is being input or terminated at any priority level, entry to it is inhibited by the setting of D46 in V25P104M5, which in turn sets the appropriate holdup bit in VOP104. This particular

hold-up covers all sorts of conditions, ranging from the complete absence of the program (indicated by a zero BA in the relevant member of V9P0M5, and no B or A "program present" bit in VOP104), through the first stages of program input (when the program, still anonymous, has a small foothold in the store), to termination by TINT A, OUT 0, OUT 1, or OUT 2.

These different conditions are indicated by the quartet V29P104M5, the various bits of each member of which having the following significance:

```
D0 = 1 during termination of a program, indicates that the termination of "OUT 8" - type procedures is awaited. The main termination subprogram is held up at one stage until the subsidiary (OUT 8) subprogram has completed sequences, which lead to this bit being cleared.
```

```
D1=0,D2=1 this combination of bits occurring during program input indicates an input holdup - due e.g. to lack of program input device, or sufficient core store - such that the input process can be terminated by TINT A. If TINT A is performed while V29P104M5 contains this pattern, it is converted to
```

```
D1=1,D2=1 which causes the input procedure to be appropriately cancelled (if TINT A is performed while V29P104M5 is non-zero, but contains some other combination in these bit positions, it is ignored).
```

```
The execution of TINT E (A or B as appropriate) while an A or B program is awaiting a program tape, i.e. while the member of V29P104M5 corresponding to that program has
```

```
D1 = 0, D2 = 1, will cause D2 to be reset to zero.
```

```
D47 = 1 during the input or termination of a program, indicates that it is a B program,
```

```
D46 = 1 OUT 1 in progress
```

```
D45 = 1 TINT A or OUT 0 in progress, or normal program input
```

```
D44 = 1 OUT 2 in progress
```

While a priority level is unoccupied, and is not involved in a program input, the relevant member of V29P104M5 contains 5 (D45 and D47 = 1) regardless of whether the priority level is A or B.

Program input is performed by the subroutines P120 and P52 - the latter handles B program input (TINT 7) in its initial stages and soon joins P120. Program input is carried out using the main subprogram (see 4.1.5) of the priority level concerned, so that effectively it is an activity having the same priority as the program it is loading. Advantage is taken of the restriction that not more than one A program and one B program can be in process of input at one time, in that only 2 sets of the parameters relevant to these activities need be stored: these

are kept in V0-2 of P120 (A programs) and P52 (B programs), as follows:

```
V0 = priority level at which input is taking place
    (negative if inactive)
```

```
V1 = program tape number (A or B)
```

```
V2 = input paper tape reader (must be cleared as soon as
    PTR is finished with)
```

The instruction sequences in P120 for A or B program input are largely identical. One subroutine, V25P104M5, which gives the priority number in VOP104 set to 4, and only if, the corresponding member of the quartet is non-zero. The various bits in each word of this quartet when non-zero, each indicate a different reason for a holdup, as follows:-

```
D43 = 1 implies a store move and/or priority change impends
D44 = " a hold-up due to the "OUT 8" subprogram
D45 = " the program is suspended by TINT 8
D46 = " the program is absent, or being input or terminated
D47 = " the main subprogram is active
```

The significance of D44 and 47 is explained in the sections dealing with subprograms (4.1.5, 4.1.6, 5.3). The termination hold-up, D46, is described in 2.5, below.

The subroutines P113 and P114 are used respectively, to set and clear bits in V25P104M5, and at the same time set or clear the hold-up bits in VOP104, as appropriate.

3. Interruption processing

3.1 Paths through the Director

As in the non-time-sharing Director (N.4) there are two main paths through the time-sharing Director from the interrupt entry at word 0 to resumption of a program. Fig. 1 shows them in outline.

The coding for both paths comprises P0 of the Director Call Program. The "standard preliminaries" mentioned in Fig. 1 are the instructions:

```
Q0; SHL+63; +=Q0;
SET B140000; =K1; ERASE;
```

which are stored in words 0 and 1 (by the Director Call Program).

Because both paths require the use of at least one Q-store, Q5 is dumped before the paths diverge. Similarly the state of the overflow and TR are recorded. In this context "recorded" implies that they are stored in the member of a quartet - V1P104M5, in this case - which pertains to the program interrupted (the priority of this program is to be found in V7P0). "dumped" implies parking temporarily in a single location.

The Short path is followed if one of the RFI's (CLOCK, FLEX, LIV, NOUV, EDT, OUT, RESET, or the abnormal absence of any RFI's at all) sends the Director down the Long path. Any other RFI (e.g. TINT, V1, V2, V3, V4, V5, V6, V7, V8, V9, V10, V11, V12, V13, V14, V15, V16, V17, V18, V19, V20, V21, V22, V23, V24, V25, V26, V27, V28, V29, V30, V31, V32, V33, V34, V35, V36, V37, V38, V39, V40, V41, V42, V43, V44, V45, V46, V47, V48, V49, V50, V51, V52, V53, V54, V55, V56, V57, V58, V59, V60, V61, V62, V63, V64, V65, V66, V67, V68, V69, V70, V71, V72, V73, V74, V75, V76, V77, V78, V79, V80, V81, V82, V83, V84, V85, V86, V87, V88, V89, V90, V91, V92, V93, V94, V95, V96, V97, V98, V99, V100, V101, V102, V103, V104, V105, V106, V107, V108, V109, V110, V111, V112, V113, V114, V115, V116, V117, V118, V119, V120, V121, V122, V123, V124, V125, V126, V127, V128, V129, V130, V131, V132, V133, V134, V135, V136, V137, V138, V139, V140, V141, V142, V143, V144, V145, V146, V147, V148, V149, V150, V151, V152, V153, V154, V155, V156, V157, V158, V159, V160, V161, V162, V163, V164, V165, V166, V167, V168, V169, V170, V171, V172, V173, V174, V175, V176, V177, V178, V179, V180, V181, V182, V183, V184, V185, V186, V187, V188, V189, V190, V191, V192, V193, V194, V195, V196, V197, V198, V199, V200, V201, V202, V203, V204, V205, V206, V207, V208, V209, V210, V211, V212, V213, V214, V215, V216, V217, V218, V219, V220, V221, V222, V223, V224, V225, V226, V227, V228, V229, V230, V231, V232, V233, V234, V235, V236, V237, V238, V239, V240, V241, V242, V243, V244, V245, V246, V247, V248, V249, V250, V251, V252, V253, V254

3.4 Timing - 20 -

The Real Time record is maintained by P60, which increases it by 1.048576 seconds every time a CLOCK RFI is detected.

The "Notional Elapsed Time" record is also maintained by P60, which adds 1.048576 seconds every time a CLOCK RFI is detected to those members of the quartet V9P104M5 which corresponds to PNU registers currently showing a peripheral hold-up.

The Run Time record for each program is recorded (in the quartet V5P104M5) in much the same way as in the non-Time-sharing Director (N.3.2); i.e. whenever a program is run in time is measured by the scaled value (including overflow) of the first subsequent CLOCK reading, which is always taken a fixed length of time after entering the Director. Before a program is resumed its run time is decreased by the scaled value (excluding overflow) of the last CLOCK reading, and also by a "path correction" which represents the sum of two time intervals from the last reading until when the path was first entered to the time of the next subsequent CLOCK reading - the latter is fixed, the former depends on the path through Director variations are introduced by the need to execute program loops both for determining the return priority stage (1 or 2 short path) and for restoring the Nest and SJNS (end of long path).

To calculate the path corrections, a similar strategem is employed to that used for timing the non-Time-sharing Director (N.3.2). A program is written which obeys OUT 3 consecutively twice, the second execution being at a different priority level. By subtracting the two run times should be negligible. The Director is suitably modified so that (a) it regards a peripheral hold-up (the sequence V9P104: K5; OR; becomes V9P104: K5; ERASE8) and (b) instead of recording real time it records the actual unscaled clock readings, using the "path corrections" to provide a correction for the path taken. The "path correction" (which is equal to 1). The sum of the difference of the two OUT 3 readings provides the program with an accurate record of the "path times" which elapsed between them, and of the number of interrupts. The path corrections can be obtained by dividing one into the other. By suitably varying the priority at which the program is run, and the number of times it has had in its Nest and SJNS when interrupted, the loop variations can be calculated as well. The Director can be made to follow any path by using the "Sense switch" positions in the RFIR (D16-21) - when any of these is non zero, the long path is followed, and one or more of them may be used to appropriate values in V14P0.

3.5 V-stores of P0
(V5,8,15,23 contain constants)
V0 = real time record, scaled V17 = LNV (non-zero if set)
V1 = 4-magn. tape unit, or 0-47 V18 = SHUP (do-it)
V6 = last clock reading, scaled V19 = EDT (do-it)
V7 = Priority of interrupted program V20 = OUT (do-it)
V9 = BA quartet V21 = *purious interrupt indicator
V13 = Time expired indicator V22 = RESET (non-zero if set)
V14 = Negative, or return priority V24=17-TAB word quartet
V16 = FLEX (non-zero if set) V28 = Director/Program marker

- 21 -

4. Hold-ups and Subprograms
4.1 Types of hold-ups - Action following EDT.
Director holding is mainly caused by the need to devise schemes for arranging internal "hold-ups" analogous to the interruptions which hold up ordinary programs. This need arises particularly in parts of the Director concerned with input/output (N.2.3).

The removal of any of these hold-ups is almost invariably signalled by an EDT interrupt. The action of the director if an EDT is therefore largely concerned with looking at all the hold-ups recorded, deciding which of them can be removed, and restarting the process which was held up.

The types of activities which have to be held up in the Time-Sharing Director fall into 5 main classes:

- Operations on W magnetic tapes. These can be classed as standard function sequences carried out by Director on several different peripherals. They are independent of priority.
- The "consolidation" within the core store of the areas belonging to the different programs occurring in the machine, and the interchanging of priorities concerned on the running of A-programs (and TINT V) - this can be regarded as a continuing process held up when it has nothing to do, and also when some of its activities are delayed by lock-outs (see 2.4).
- "Queued" operations. When several programs, possibly including the Director, wish to make use of the same peripheral unit, it is natural to deal with them on a "first come, first served" basis - i.e. each request for a transfer on this unit is placed in a queue. As each transfer ends, the corresponding item from the queue is removed, the queue moves up one place, and the next transfer is started. These operations are, by their nature, independent of priority. A new transfer request is placed at the end of the queue and a subsequent cannot change places with any other item in the queue.
- The initiation of A-program input. This does not refer to the detailed business of program input, which is handled by subprograms but the Director's overall task of keeping the machine supplied with A-programs. This may be regarded, like (b), as a continuing process which is held up at times when there is no suitable priority available.
- Operations within the Director which have to be given the same relative priority (in respect of the order in which hold-ups are removed) as the programs on whose behalf these operations are carried out. These operations include the allocation of storage for program input, the allocation of peripheral units for program input or to programs already input, as well as input-output operations carried out by the Director on behalf of a program.

- 22 -

The subroutine P5, which is entered whenever V19P0 (the EDT marker) is found to be non-zero, takes account of all these possibilities by examining in turn the recorded hold-ups within all these classes.

The appropriate subroutines, in order of entry, are:-

- P100 (deals with W magnetic tapes)
- P105 (consolidates store, etc.)
- P59 (deal with Flexowriter queue)
- P155 (deal with OUT 8 queue)
- etc.
- P120 (initiate A-program input)
- P11 (deal with subprograms).

The order in which these subroutines are entered is important, for the following reasons:-
- P105 must precede P11, as the change in status of a magnetic tape unit from W to L by P100 may remove a holdup from a subprogram, which will be detected by P11.

- P105 must precede all the queue operations, because some of them may be held up awaiting store moves or priority changes (see reference to P103 in 2.4). For the same reason P105 has to precede P11. P105 should also precede P120, because the initiation of the input of a new A-program may be made possible by the actions of P105.

- the queue operations must precede P11, because the end of a queued transfer may remove a subprogram hold-up.

- P120 should precede P11, since the former sets up a subprogram to be re-activated by the latter.

These subroutines are described in detail below.

After all these subroutines have been entered P5 ends by carrying out the examination of the PNU store prescribed in 1.3. Should this lead to the clearing of the part of the store which contains the set of the appropriate members of V13P104M5, and a predetermine return priority number may be left in V14P0 - the coding should be consulted for the algorithm used.

4.1.1 Dealing with W Magnetic Tapes (P100)
This process, which is almost identical to that used in the non-Time Sharing Director (N.6.4), ensures that any magnetic tape unit assuming status W undergoes a sequence of operations which are carried out by one or more of a number of "bricks". Each magnetic tape unit has its own brick indicator (B.I.) which indicates which bricks remain to be dealt with. The B.I. is cleared when the unit is removed from the status of unit W to something else. The value of the B.I. is specified when the unit is given status W, and the specified bricks are then obeyed in numerical sequence. The bricks are numbered 1-6 and carry out the following functions:-

Brick 1 - Count blocks back to BT.
Brick 2 - Skip back to BT without counting.
Brick 3 - Read label block.
Brick 4 - Alter status from W to U, typing full details.
Brick 5 - As brick 4 changing status from W to L.
Brick 6 - Alter status from W to L, typing nothing.

- 23 -

Entry to a brick which occurs whenever the unit concerned is not busy, is arranged by P100, which provides the brick with the address of the first record in the unit number in C6 and M6. No brick is permitted to alter O8. The entry points to bricks 1-6 are labels 1-6, respectively, of P101.

A brick must be entered over and over again until its job is complete - this is decided by the brick itself, which returns control to P100 at one re-entry point (SP100), when it is not completed, another (6P101) when it is - in the latter case P100 updates the B.I. to ensure that the right brick is entered next. The B.I. for unit n, which is stored in D0-23 of V(n-6)P100, contains the floating point number

$$2^{-8} + 2^{-9} + 2^{-10} + \dots$$

where a, b, c, ... are the numbers of the uncompleted bricks.

The label block of the tape on unit n is read into the three words which begin at V(3n-18)P101.

To alter the status of a magnetic tape unit to W, it is necessary to obey

J510P100;

with the unit number in M6, and the non-normalised brick indicator in N1 - in this form D33=1 indicates brick 1 is to be obeyed, D34=1 indicates brick 2, D35=1 brick 3, etc.

D24 of the B.I. word holds a Parity indicator, and D25-47 the block count used by brick 1.

4.1.2 Action of P105

Some mention of the actions involved in store consolidation and priority interchange was made in 2.4. These actions are dealt with by P105, which performs the following steps:-

(1) Examines and clears a marker (V9P105), and skips step (2) if it was zero.

(2) For each priority in turn, effects any specified store movement (but only if the area involved is not locked out) and any specified priority interchange (but only if neither of the areas assigned to the specified priorities are locked out).

"Store movement" involves not only moving the program in the store and altering its RA (in V9P0MS) but also modifying any absolute addresses referring to locations within the "moved area" which may be stored as parameters in any of the "Queues" (see 4.1.3) or any of the subprogram hold-up parameter stores (see 4.1.5).

"Priority interchange" involves the following:-

- interchanging bits of V9P104.

- 24 -

- (with additional complications) in V9P0MS and V24P0MS.

- altering, if need be, the priority numbers recorded in V7P0, V9P52 and V9P120, and similar locations.

- interchanging the sets of subprogram parameters for the pairs of subprograms pertaining to the priorities involved (see 4.1.5).

(3) Re-sets up all the priority interchanges required by the A-program concept, regarding only those & priorities which are neither occupied by B programs, nor already committed to inter-change with B, or B-occupied, priorities, as eligible for interchange. Any new priority interchange set up in this step also causes V9P105 to be set non-zero.

(4) If there are gaps in the store, either between two A-programs or between the top of the store and the nearest A-program, sets up the store move which will close up the gap nearest the top of the store. Any new move set up causes V9P105 to be made non-zero.

(5) Carries out a similar service for B-programs as (4) does for A-programs, relative to the bottom of the store.

(6) If V9P105 is now non-zero, returns to step (2).

(7) For each priority, causes the hold-up bit D43 of the corresponding member of V9P104 to be set or cleared according to whether or not the priority is still an outstanding store move or priority interchange at that priority. If there are any outstanding, V9P105 is left non-zero.

4.1.3 Queued transfers

The format of transfers in any queue (e.g. the Flexowriter queue, or with additional complications) in V9P0MS and V24P0MS.

(a) Peripheral transfers - the final absolute address of the store, and the initial absolute address of the peripheral device or a locked-out area of stores, or zero.

(b) the subprogram is waiting for the end of a "queued" peripheral transfer, or

(c) the subprogram is waiting to be allocated a magnetic tape with a specified identifier, or

(d) the subprogram is waiting to be allocated some other type of peripheral unit.

The total subprograms are numbered 0-9. Subprograms 0 and 1 belong to the Director - the former is used by P4 for typing out the "TINT" query, the latter is used by OUT 8 and TINT M. The remaining eight are allocated to the four priorities; thus, subprograms 2-9 are allocated to priority 0, 10-17 to priority 1, etc.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the same as the contents of the subprogram which was entered before it. If they are, then the subprogram is cleared, and the next subprogram is entered.

When a subprogram is entered, it is checked to see if it is a "queued" transfer. If it is, the contents of the subprogram are examined to see if they are the

P10 has several entry points, of which one (the "normal" one, JSPI10) is conditional, in the sense that the subprogram is only suspended if a peripheral holdup is present: all the other entries are guaranteed to suspend the subprogram. The following requirements apply to all entries:

- (i) the correct parameters must be in Q4, and when the subprogram is entered, number 2, 4, or 8, the correct priority number must be in M5 (to address V25P104M5, etc.).
- (ii) entry must be made using a "jump to subroutine" instruction. If the subprogram is actually suspended, the link thus planted will be stored in SHUL, so that when the subprogram is re-entered (and this applies even if it is not suspended) it will be at the instruction immediately after the "jump to subroutine".
- (iii) a parameter must be supplied which states how many items (up to 32) are to be transferred from SJNS to SHUL* if this entry causes the subprogram to be suspended.
- (iv) if the subprogram is suspended, P10 is left (after having deposited the appropriate items in SHUL, SHUL* and SHUP) by obeying EXIT 1 or EXIT 2, and a further parameter must be supplied to say which of the two exits is obeyed; leave the top two go to SHUL*, and the third is used when EXIT 2 is obeyed. The link planted on entry to P10 does not concern its position to SJNS, where the program is re-entered subsequently by P11; the SJNS will contain (starting at the top and working down): the 2 items from SHUL, a link to return control to P11, and links to return control from P11 to P5 and from P5 to the long path.

Note that if a subprogram enters P10 but is not suspended, it will continue at the next instruction after the "jump to subroutine" with the SJNS, SHUL, etc., unchanged.

The use of the SJNS to determine where control shall be transferred when P10 is left ensures that P10 can be used to initiate a subprogram, i.e. to suspend an activity in such a way that it can be re-entered as a subprogram by P11; provided that this initiation, or initial suspension, takes place within a subroutine, so that there is in SJNS a link which may be used, by EXIT 1 or EXIT 2, when transferring control from P10 after the latter has set up SHUL, SHUL*, and SHUP. Because there may be some significance in whether this link is used by EXIT 1 or EXIT 2, the choice must be specified in the link parameter.

Once the subprogram has been initiated, and re-entered by P11, any subsequent suspension by P10 will cause control to revert to P11, via the link planted by the latter - i.e. the subprogram behaves like a subroutine of P11. Whether EXIT 1 or EXIT 2 is used to return control to P11 is immaterial, as either will reenter P11 at the point where it updates Q4 and goes on to examine the next subprogram.

Whenever P10 suspends subprogram 2, 4, 6 or 8, it also sets Q47 of the appropriate marker of V25P104M5 so that during its suspension, the priority of which this is the main subprogram is also held up.

P10 is not used to close down a subprogram, because, as explained in 4.1.5, when a subprogram is re-entered its SHUL is cleared. All that a subprogram needs to do to close itself down is to obey EXIT 1 or EXIT 2 using the link planted by P11.

The full power of the P11-P10 combination will be shown in the examples of its uses in connection with FLEX and OUT interrupts.

Entries to P10 are as follows:-

- (a) JSPI10: obeyed with the link parameter in N1. The subprogram is suspended if, and only if, the store area defined by the absolute address D0 of N1 is still locked-out, i.e. C7 is busy. The suspension will be the first type of subprogram hold-up, i.e. D0 of SHUL is made zero, and Q7 is transferred to SHUP, by P10. When the subprogram resumes, at the instruction after JSPI10, the link parameter will have been erased from N1 and Q7 will be in status quo.
 - (b) JS2P10: obeyed with the link parameter in N2. N2 contains a quantity which P10 will OR with the link before transferring the latter to SHUL. N1 contains the quantity to be transferred by P10 to SHUP. This operation is unconditional. P10 will erase N1, N2, D1, where the subprogram is suspended, or P11 (at the next instruction following JS2P10) the disposition of items in Q7, the Nest, etc., will, of course depend on the type of suspension determined by D0 of SHUL, and SHUP (see 4.1.5).
 - (c) JS3P10: obeyed with the link parameter in N1. This entry is exactly equivalent to obeying ZERO; Q7; JS2P10; (i.e. like JSPI10, with suspension guaranteed).
 - (d) JS4P10: exactly equivalent to obeying ZERO; ZERO; ZERO; JS2P10; this is a public method of suspending, or re-entering, a subprogram so that a re-entry next time P11 is entered is guaranteed, since D0 of SHUL=0 and SHUP=0. SHUL* is also cleared.
- Apart from the link parameter, the entry JSPI10 is the almost exact equivalent of the JS/RWLS/ used in the non-Time-Sharing Director (N6.2).
- Entry 2 (JS2P10) is so flexible, permitting any type of suspension to be set in SHUL and SHUP, that it is used to set up almost all the subprogram hold-ups implicit in other subroutines, e.g. P8 and P58. P8, the subroutine to set typewriter query in the flexowriter queue (4.1.3), requires the link parameter to be P11 (as the priority is last), after inserting the transfer parameters in the queue, i.e. J2B10... P58 is subroutine used by OUT 4 and OUT 10 which, given a 2-word identifier in N1

and N2, will suspend the subprogram in the correct manner if the required tape is not available. To do this it will obey:

```
    RVN; (interchange 2 words of identifier)
    =E204M; (store 2nd word in SHUP*)
    SET1; (link parameter - to preserve link for P58 in SHUL*)
    ZERO; NOT; STR; NEW; (sets D0 in N1)
    CAD; (1st word of identifier to N1, for SHUP)
    JS2P10; (to hold up subprogram - SHUL has D0=1)
```

When the subprogram is reentered by P11, which will occur when the new identifier marker, VOP101, is non-zero, the identifier will be transferred by P11 from SHUL and SHUP* to N1 and N2. The link for P58 will be in SJNS, so that the position is exactly as it was when P58 was first entered. The next instruction is therefore: JR98;

Note that a subprogram may address its SHUL, SHUP, SHUP*, SHUL*, as E04M, E104M, E204M and E304M, respectively.

4.2 FLEX. The action of P4

The subroutine P4 is entered from the long path if FLEX has occurred, i.e. if V16P0 is found to be non-zero. Its first action is to set V19P0 (the EDTR marker) non-zero, to ensure that by pressing the Interrupt key on the Flexowriter the operator can force an "artificial" EDTR.

Subprogram P0 will be used to type the query "TINT". If this subprogram is found to be active (i.e. its SHUL word, V1P11, is non-zero) P4 exits to the long path without clearing V16P0 - this will cause the Director to cycle round the inner loop of the long path until the FLEX RFI can be dealt with.

If V1P11 is zero, P4 sets up Q4 with the correct parameters for subprogram 0 (a copy of V0P11) and enters P8 to put the correct parameters for typical programs into the paper tape puncher. It then sets up a parameter of zero. As explained in 4.1.6, this will cause SHUL (with D0=0) and SHUP (with D0=1) for subprogram 0 to be set up, and will then cause control to be returned to the long path when P8(P10) obeys the link (using EXIT 1) which was planted by the entry to P4. The subsequent actions of P4 therefore take place in subprogram 0, following re-entry by P11 after the query has been answered (giving EDTR automatically) and checked.

The reply to the query is checked to ensure it is a letter in the right range. If it is not, the query is typed again (preceded by AGAIN); this time P8 transfers control back to P11 rather than to the long path.

With a satisfied answer, P4 re-reverses and then prints a blank depending on the value of the letter typed. This branch is made to look like a subroutine jump, so that in effect each TINT is dealt with by a subroutine of P4.

These subroutines are:

```
    TINT A - P33
    " B - P44      TINT M - P45
    " C - P37      " P - P48
    " G - P39      " R - P50
    " H - P40      " S - P51
    " I - P41      " T - P52
    " J - P42      " U - P53
    " L - P44      " V - P54
```

and when each is entered in this context, there will be four links in SJNS: the links to return control to P4, from P4 to P11, from P11 to P5, and from P5 to the long path.

Each of these subroutines has two exits. EXIT 2 is the normal exit for a successful execution. EXIT 1 is obeyed when any check fails and causes the TINT query (preceded by AGAIN) to be repeated.

To illustrate the ways in which the subroutines P10, P11, P4 and the subroutines P12, P13, P14, P15, and P16 are discussed below, some detail. The other TINT subroutines are quite straightforward, and examination of any of them will show that they work in a simple way, and finish by obeying either EXIT 1 (failure) or EXIT 2 (successful execution). Following the successful execution of any of these subroutines, Q4 is liable to have been altered, so the instructions to which EXIT 2 leads are:

```
V0P11; = Q4; EXIT 1;
```

which sets Q4 to the value required by P11, and then return control to P11, leaving subprogram 0 closed down. The value of M5 is immaterial when P11 is resumed.

The subroutines P31 and P56 are entered to get the priority number of the program concerned into M5. V29P104M5 is examined to see whether the program is in process of termination (V25P104M5 nonzero) or not.

If the program is already being terminated, P33 merely sets a 1 in D1 of V29P104M5 if D2 of that word is nonzero (see 4.1.5), and then ends EXIT 2.

If the program is not already in process of termination, TINT A must end it. To do this, it first clears D45 of V25P104M5 (see 2.4.1) and then, by obeying JS2P10 with a link parameter of -1 (which causes P10 to clear S0), and end by obeying EXIT 2) sets up the main subprogram of the priority to be terminated, closing down subprogram 0 as shown below.

The steps performed by P33 in doing this are:

```
JSP20; (this most useful subroutine sets up in Q4
    the parameters appropriate to the main subprogram of the
    priority whose number is given in M5)
    ZERO; NOT; (link parameter)
    ZERO; (to ensure D0 = SHUL=0)
    C7; (sets the first following the program letter; if it is +
        the TINT A procedure is different. Here it is
        masquerading as the final absolute address (in the range
        0-63) of a potentially locked-out area, and will thus be
        stored in SHUP. Since locations 0-63 are never locked-
```

out this is more or less legitimate).

JS2P10; (this causes SHUL and SHUP of the priority's main subprogram to be set up as indicated by the parameters in N1 and N2. SHUL* is cleared and P10 obeys EXIT 2 to return control to P4. P4 now restores Q4 to its original value, for subprogram 0, and opens the link to return control to P11 which will soon come to deal with the subprogram, namely the main subprogram 0 is left closed down since its SHUL is still zero.

When the new subprogram is re-entered by P11, M7 will not contain the character following the program letter. If this is + (octal 35), an ending number(4) is placed in V0P29 and P2 entered at reference 2 (see 5.1), causing the Nest and SJNS to be typed out before entry to P15 which sets off the normal termination procedure. Otherwise P15 is entered direct.

Note that TINT A does not check whether the priority's main subprogram is suspended or closed down, before taking it over. This characteristic is shared by TINT I (see 4.2.2). It implies that any activity carried out by the main subprogram while its priority is still busied (i.e. when V0P10 is zero) will be suspended if it is re-entered by P11.

The second part of TINT A is to suspend itself if a paper tape read error occurs. This is done by obeying EXIT 2, which is the same as the failure indicator required by /RWB/ in the non-Time-Sharing Director - N7.1. The link parameter is unconditional, and the priority must contain a non-zero "Ending number"; otherwise the familiar "REACT" - query will be typed later, and termination only occurs if the value of M5 is actually -2.

Another interesting point is that, once the subprogram is re-entered, it may return to the same setting-up point (in P52 this is: SET; SET-2; ZERO; JSPI13A); to suspend itself if a paper tape read error occurs. The same link parameter is used because the P52-P10 link must still be preserved in SHUL*, and it does not matter if P10 returns control to P11 using EXIT 2.

Note that when the subprogram suspends itself because there is no paper tape reader available, it does not use the convention of setting D0 of SHUL = 1, and SHUP. D0-type number, because this does not give the operator any opportunity of using TINT A to terminate the input request (hence the use of P14). If TINT A occurs, the subprogram clears V29P2 and enters P12 (the Program Failure routine) directly.

Until falling into the trap mentioned in 4.1.3, the first action of P10 is to clear V20P0, the OUT marker. Then it adds 3 syllables to the programs Nest and SJNS, and checks that there is an "OUT number" in the right range stored at the top of the Nest, or V0P1. A copy of this is transferred to the SJNS, and this is part of the subprogram, other than the routines outside it. Also the link parameter must be negative since EXIT 2 is required: so its value is actually -2.

It will be noted that, in these four examples, some other subprogram, besides P10, is involved. However the other subprogram is never set up, and the subroutines P12, P13, P14, P15, and P16 are discussed below, some detail.

The other TINT subroutines are quite straightforward, and examination of any of them will show that they work in a simple way, and finish by obeying either EXIT 1 (failure) or EXIT 2 (successful execution). Following the successful execution of any of these subroutines, Q4 is liable to have been altered, so the instructions to which EXIT 2 leads are:

```
V0P11; = Q4; EXIT 1;
```

which sets Q4 to the value required by P11, and then return control to P11, leaving subprogram 0 closed down. The value of M5 is immaterial when P11 is resumed.

The subroutines P31 and P56 are entered to get the priority number of the program concerned into M5. V29P104M5 is examined to see whether the program is in process of termination (V25P104M5 nonzero) or not.

If the program is already being terminated, P33 merely sets a 1 in D1 of V29P104M5 if D2 of that word is nonzero (see 4.1.5), and then ends EXIT 2.

If the program is not already in process of termination, TINT A must end it. To do this, it first clears D45 of V25P104M5 (see 2.4.1) and then, by obeying JS2P10 with a link parameter of -1 (which causes P10 to clear S0), and end by obeying EXIT 2) sets up the main subprogram of the priority to be terminated, closing down subprogram 0 as shown below.

The steps performed by P33 in doing this are:

```
JSP20; (this most useful subroutine sets up in Q4
    the parameters appropriate to the main subprogram of the
    priority whose number is given in M5)
    ZERO; NOT; (link parameter)
    ZERO; (to ensure D0 = SHUL=0)
    C7; (sets the first following the program letter; if it is +
        the TINT A procedure is different. Here it is
        masquerading as the final absolute address (in the range
        0-63) of a potentially locked-out area, and will thus be
        stored in SHUP. Since locations 0-63 are never locked-
```

out this is more or less legitimate).

JS2P10; (this causes SHUL and SHUP of the priority's main subprogram to be set up as indicated by the parameters in N1 and N2. SHUL* is cleared and P10 obeys EXIT 2 to return control to P4. P4 now restores Q4 to its original value, for subprogram 0, and opens the link to return control to P11 which will soon come to deal with the subprogram, namely the main subprogram 0 is left closed down since its SHUL is still zero.

When the new subprogram is re-entered by P11, M7 will not contain the character following the program letter. If this is + (octal 35), an ending number(4) is placed in V0P29 and P2 entered at reference 2 (see 5.1), causing the Nest and SJNS to be typed out before entry to P15 which sets off the normal termination procedure. Otherwise P15 is entered direct.

Note that TINT A does not check whether the priority's main subprogram is suspended or closed down, before taking it over. This characteristic is shared by TINT I (see 4.2.2). It implies that any activity carried out by the main subprogram while its priority is still busied (i.e. when V0P10 is zero) will be suspended if it is re-entered by P11.

The second part of TINT A is to suspend itself if a paper tape read error occurs. This is done by obeying EXIT 2, which is the same as the failure indicator required by /RWB/ in the non-Time-Sharing Director - N7.1. The link parameter is unconditional, and the priority must contain a non-zero "Ending number"; otherwise the familiar "REACT" - query will be typed later, and termination only occurs if the value of M5 is actually -2.

Another interesting point is that, once the subprogram is re-entered, it may return to the same setting-up point (in P52 this is: SET; SET-2; ZERO; JSPI13A); to suspend itself if a paper tape read error occurs. The same link parameter is used because the P52-P10 link must still be preserved in SHUL*, and it does not matter if P10 returns control to P11 using EXIT 2.

Note that when the subprogram suspends itself because there is no paper tape reader available, it does not use the convention of setting D0 of SHUL = 1, and SHUP. D0-type number, because this does not give the operator any opportunity of using TINT A to terminate the input request (hence the use of P14). If TINT A occurs, the subprogram clears V29P2 and enters P12 (the Program Failure routine) directly.

Until falling into the trap mentioned in 4.1.3, the first action of P10 is to clear V20P0, the OUT marker. Then it adds 3 syllables to the programs Nest and SJNS, and checks that there is an "OUT number" in the right range stored at the top of the Nest, or V0P1. A copy of this is transferred to the SJNS, and this is part of the subprogram, other than the routines outside it. Also the link parameter must be negative since EXIT 2 is required: so its value is actually -2.

It will be noted that, in these four examples, some other subprogram, besides P10, is involved. However the other subprogram is never set up, and the subroutines P12, P13, P14, P15, and P16 are discussed below, some detail.

The other TINT subroutines are quite straightforward, and examination of any of them will show that they work in a simple way, and finish by obeying either EXIT 1 (failure) or EXIT 2 (successful execution). Following the successful execution of any of these subroutines, Q4 is liable to have been altered, so the instructions to which EXIT 2 leads are:

```
V0P11; = Q4; EXIT 1;
```

which sets Q4 to the value required by P11, and then return control to P11, leaving subprogram 0 closed down. The value of M5 is immaterial when P11 is resumed.

The subroutines P31 and P56 are entered to get the priority number of the program concerned into M5. V29P104M5 is examined to see whether the program is in process of termination (V25P104M5 nonzero) or not.

If the program is already being terminated, P33 merely sets a 1 in D1 of V29P104M5 if D2 of that word is nonzero (see 4.1.5), and then ends EXIT 2.

If the program is not already in process of termination, TINT A must end it. To do this, it first clears D45 of V25P104M5 (see 2.4.1) and then, by obeying JS2P10 with a link parameter of -1 (which causes P10 to clear S0), and end by obeying EXIT 2) sets up the main subprogram of the priority to be terminated, closing down subprogram 0 as shown below.

The steps performed by P33 in doing this are:

5.2.2 P16 (OUT 1)

After various checks, the value 2 (A-program) or 3 is set in V29P104M5, and D46 of V25P104M5 made non-zero. The main subprogram is suspended if the store area of the program is locked out, if the second (OUT 8) subprogram is not closed down, or if VOP120 (A program)/VOP52 (B program) is positive, and will not be resumed until all these conditions are absent. On resumption, the priority number will be stored in VOP120 or VOP52, to indicate that A- or B-program input is going on.

Before this suspension a message "OUT 1" is typed: the operator will be aware that the suspension has not been lifted if the new program identifier is not typed subsequently.

Finally P16 obeys J2P120 to transfer control to the program input routines. Parameters for P120 will have been left in V1P104M5 (the amount of store used by the previous section, not to be exceeded by the new one) and in V45P104M5 (the identifier of the required new program section).

Any failures after entry to P120 will of course be dealt with by P121 (CRNP failure), which recognises that the program input was initiated by OUT 1 (since D46 of V29P104M5 is non-zero) and accordingly transfers control to P2 - having first set the ending number 5 in VOP2 to guarantee termination.

5.2.3 P18 (OUT 3) and P24 (OUT9)

These two very simple routines are examples of the way in which OUT routines may, until their first suspension, rely on the program's Nest staying in the V-stores of P1. The appropriate times (Run or Real) are transferred to VOP1 (N1 of the program's Nest). Thus the OUT number is erased and the result inserted with just one instruction.

5.2.4 P19 (OUT 4) and P61 (OUT 10)

These routines both, after performing all the necessary checks, use P58, whose action has been described in 4.1.6. If the required identifier is one-word, an artificial second word is introduced which has D0=2, D2=47 = all 1's.

A marker is set in V45P104M5 to distinguish OUT 4 (marker non zero) from OUT 10 (marker zero).

5.2.5 P20 (OUT 5) and P21 (OUT 6), P26

The behaviour of the OUT 5 routine depends largely on whether the program obeying OUT 5 is A or B. An A program can only be allocated to a unit which is free, and if there are none available, if there are none available, a B program will be allocated a P unit if there is one (a unit may be pre-allocated to a B program either as the result of TINT T, which can pre-allocate the input paper

- 41 -

tape reader, or following the de-allocation of a unit previously allocated by OUT 5*) if there is a free unit of the required type, a unit may be allocated and if there is none of these the "AWAITC TYPE 4" message is typed and the subprogram suspended (with SHUL negative, and SHUP = type number + D0). This suspension is cleared either by a unit of the right type becoming available, or by TINT A or I0.

OUT 5* (when the type number supplied by the program in N2 has D44 added) is dealt with as follows: When a B-program asks for a unit of a certain type in this way, P20 inserts a marker bit in V21P104M5. The bits of this word are assigned to type numbers in such a way that type OUT 5 is assigned by bits 8a - 8a b). When a B-program sees OUT 5 to re-allocate a unit whose type bit in V21P104M5 is non-zero, P21 clears that bit and the unit is given status P, rather than U as would normally happen. If a B-program tries to use OUT 5 or 5* to allocate a unit, specifying a type whose bit is non-zero, Program Failure 2 occurs even if a unit of the required type is available. The system thus has exactly the same effect as in the non-Time-Sharing Director.

This only applies to B-programs. Because any unit (not magnetic tape) which is deallocated by a P-program is given status U, and a unit of any other type is pre-allocated, OUT 5 and OUT 5* are identical for A-programs. For an A-program V21P104M5 instead contains a record of all the units pre-allocated to the program: D47 = unit 0, D46 = unit 1, D37 = unit (octal) 12, etc.

This correspondence of bits with unit numbers, which is the same as that used in CPDAR (N2.1.4), also applies to the copy of the programs CPDAR kept in V17P104M5, which is updated by P20 and P21 as units are allocated and de-allocated.

When de-allocating a magnetic tape unit, P21 sets a brick indicator and descriptor (see 4.1.1), unless the unit's status is N. For any other type of unit, P21 has to wait until the unit is not busy before changing its status to P or U. Therefore if the unit is busy the subprogram is suspended, with SHUL positive and the unit number in D0-15 of SHUP. The end of such a suspension will not usually be signalled by an EDT and so may have to wait for an EDT from some other source. As has already been pointed out (5.2.1) if TINT A is used to terminate the program while this suspension is still present, the unit, if still busy, may have its transfer forcibly terminated.

The "descriptor" of unit n is stored in Vn of P26. Each descriptor is laid out as follows:

```
D06 - 11 = (octal) 0207
D12 - 23 = Type number (octal) in character form
D24 - 29 = Status character
D30 - 41 = Unit number (octal) in character form
D42 - 47 = program letter if any
```

- 42 -

P26 is used for finding a unit or units having a particular type, status and program letter. It has 3 entries and 2 exits. The first 2 entries (JS1P26, JS2P26) require in N1 a parameter comprising:-

```
D33 - 38 = required type number
D39 - 44 = required status character
D45 - 47 = 4 if program letter is to be blank
          = 0, 1, 2, 3 if P, Q, R, S.
```

The list is searched backwards, i.e. through decreasing unit numbers. The first entry (JS1P26) is used if the search is to start at the highest unit number, the second entry (JS2P26) if the search is to resume after the last unit found (i.e. with Q6 unaltered since the last EXIT 2). EXIT 1 is obeyed if no unit was found: the parameter in N1 will have been erased. EXIT 2 is obeyed after a successful search, when C6 and M6 will both contain the number of the required unit, with I6 = -1. After EXIT 2, N1 will contain a modified version of the original parameter, and if desired the search may immediately be resumed, without re-specifying the original parameter, by obeying JS4P26.

(Compare P26 with /RPS17 and /RPS27 of the non-Time-Sharing Director - N5.3.).

5.3 P23 and OUT 8

All transfers executed by the Director on behalf of a program (OUT 8 is the obvious example) have certain features in common. A system has therefore been evolved into which additional facilities of this type can be fitted without too much difficulty, and the framework of this system forms part of P23, the OUT 8 routine.

Besides using the main (even) subprogram, the system makes use of the second (odd) subprogram of the priority instigating the transfer - i.e. for priority p, subprograms 2p + 2 and 2p + 3. The second subprogram is used because these transfers, once started, do not require the program to be held up - the transfer area being protected by ordinary lock-outs - but do require that when they end, a procedure be used which not only checks the transfer but also signals that a further transfer may now take place. For this procedure the main subprogram is not suitable, since it would hold up the program unnecessarily, but the second subprogram is ideal.

All Director transfers of the OUT 8 type pass through 4 stages. The first stage involves erasing the OUT number from the program's Nest and checking that the parameters provided by the program in its Nest are valid - e.g. that the addresses are reasonable. It is assumed that OUTS which perform this sort of transfer will supply all parameters in N2, and that these will implicitly or explicitly specify the relative addresses of the transfer area.

This stage will end after these parameters have been copied into V45P104M5 (and possibly V49P104M5) without erasing them from the program's Nest, and the initial and final absolute addresses of the transfer area into I7 and M7. P13 (q.v.) is a useful subroutine for this purpose - it also checks that these addresses are valid. The first stage ends by obeying JS1P23.

- 43 -

The subroutine P139 causes the main subprogram to be suspended until the area involved in the transfer is not locked out, and the second subprogram is closed down (E1M4 = 0). When both these conditions are satisfied, the second subprogram starts to transfer the data. It does this through the main subprogram is able to check, and if necessary alter, the transfer area itself; and, because the second subprogram has been closed down, it may check any failure indicators (e.g. parity) left by the previous transfer of this type, which must now be over. It is a feature of the Time-Sharing Director that each priority may only carry out one transfer of this type at a time, because each transfer requires the use of the second subprogram.

Causes of program failure (JP2) may be found in the second as well as in the first stage. At the end of the second stage the main subprogram is ready to initiate the second subprogram which is going to actually perform the transfer. To do this it calls a subprogram called J37P23 with certain parameters in the nesting store. They are:-

```
N1 - SHUP* of second subprogram. This is to contain, in
      D16-31 and D32-47, the initial and final relative
      addresses of the transfer. D0-15 may contain any
      other parameter.
```

```
N2 - SHUL of second subprogram, defining its point of
      entry.
```

```
N3 - a marker which is zero if the transfer is going to
      be queued, nonzero otherwise.
```

On entry to reference 97, P23 transfers N1 and N2 to SHUP* and SHUL of the second subprogram (E21M4 and E1M4), and clears SHUL* (E31M4) and SHUP (E11M4) - the latter guarantees that at the very next opportunity, i.e. as soon as the main subprogram returns control to P11, the second subprogram will be entered. The main subprogram now obeys JS3P29 to erase the parameter word (originally in N2) from the program's Nest, and returns control to P11 - it does this by obeying EXIT 1 (thus closing down the main subprogram altogether) if the transfer is not to be queued, in which case D44 of V25P104M5 is first made non-zero, or by obeying JS4P10, if it is queued.

Stage three of the process takes place in the second subprogram. This will extract the parameters from SHUP*, make the addresses absolute, and initiate the transfer, either by entering it in a queue or by using an ordinary transfer instruction. If the latter, D44 of V25P104M5 is cleared as soon as the transfer is started, this will cause lock-outs to be set up on the transfer area, making it safe to allow the program to resume; the second subprogram is suspended until the end of the transfer, and the first subprogram is already closed down. However, if the transfer is queued there is no guarantee that it will start initially, and it is not safe to return control to the program. It will not be possible for the second subprogram to hold up the program only until the lock-out is set up, because it will itself probably be suspended until the end of the transfer. This is why, in this case, the main subprogram is not closed down; it remains active (thereby holding up the program) and is still running.

- 44 -

only closes itself down when either the area becomes locked out, or the transfer parameters (taken from E21M4, which must not be altered by the second subprogram) are found to be absent from all the peripheral queues. The main subprogram repeatedly tests these conditions, suspending itself with JS4P10 after each test cycle, until at last it obeys EXIT 1 when it is safe to return to the program. It is possible that a TINT A or TINT I0 may close down the main subprogram before this time; in these special circumstances this does not matter, and the second subprogram is not affected.

The second subprogram, continuing the third stage of the process, must remain active until the transfer is complete and checked (after repetition if need be). If it is necessary for this subprogram to hold up the program it can reset D44 of V25P104M5. The parameters in its SHUP* must not be disturbed if the transfer is queued. Even if no check on the working of the transfer is needed at its end, the ending of the third stage must be delayed until the transfer is complete; so that, if the program is terminated before the transfer ended, it will still be possible to indicate definitely when all lock-outs on the program's store have been removed.

The fourth and final stage of the process is entered by obeying J39P23. At this point (at which P15 will restart the second subprogram if it finds the latter already closed down see 5.2.1) V29P104M5 will be examined: if it is positive, the second subprogram can be closed down, but if it is negative, indicating that the program is being terminated, a winding-up procedure for all OUT 8 type activities must first be performed, culminating in D0 of V29P104M5 being made zero again.

Before the second subprogram is closed down, D44 of V25P104M5 is cleared. Also, since the main subprogram may be suspended in P139 waiting for the closing-down of the second subprogram, the subprogram parameters in Q4 are adjusted as if they were referred to the basic subprogram. In this way, when EXIT 1 is obeyed to close down the main subprogram and return to P11, the latter will now try to re-enter the main subprogram. In other words the subprogram counter is effectively stepped back one place instead of forward when the second subprogram is closed down.

The common features of all these OUT - instigated peripheral transfers are thus:

- a first stage, leading to P139
- a second stage leading to 97P23
- (two stages are carried out by the main subprogram)
- a third stage carried out by the second subprogram, which performs the transfer and leads to
- the final stage, entered at 99P23, leading to the closing down of the second subprogram (and reverting to the main subprogram)

to the main subprogram)

To fully appreciate the way these common features are exploited, it is worth examining and comparing the coding for the different types of OUT 8 process in P23, and for the OUT 11 - OUT 12 (drum transfer) routines P151 and P152.

- 45 -

6. Subroutine Index

The table below gives against each subroutine a brief title, the number of the section (if any) of the preceding text in which the subroutine is described most fully, and, where applicable, an indication of which Q-stores it alters. A single asterisk against a subroutine indicates that Q4 is required to preset with subprogram parameters.

P0 Long and Short Paths (3.1)
P1 Dump programs Nest, etc. (3.3)
**P2 Program Failure (5.1)
**P3 OUT (5)
P4 FLEX (4.2)
P5 EDT (4.1)
P6 Enter Message words in Flexowriter queue (4.1.3)
P7 Alters Q7, C5
Alters parameters in Flexowriter queue (4.1.3)
Alters Q7, C5

**P8 Flexowriter query (4.1.3)
P9 Convert cell count, used by P12
**P10 Suspend subprogram (4.1.6)
P11 Re-enter subprograms in order (4.1.5)
P12 Type program's Nest, etc. (5.1)
Alters Q6, Q7

Set and clear "software HU" bit in VOP104
Alters C5

**P13, P14 Set or clear given bit in V25P104M5 (2.4)
P15 Set priority related parameters in C6 and M6
Alters C6, M6

P16 Form table of BA/NOL words in order of decreasing BA
Alters Q6, Q7

P17 Form parameters of the "Hole in the Middle"
Alters Q6, Q7

Set up parameters before program input (4.2.4)
Alters Q6, Q7

**P18 Set up main subprogram parameters in Q4 (4.2.1)
P19 Fetch character to 07 (4.2)
P20 Fetch number (4.2)
P21 Fetch letter if any (4.2.1)
P22 Check tape identifier
Alters C6

P23 Transfer area to be queued (4.2.1)
P24 Set up transfer area to be queued (4.2.1)
P25 Set up transfer area to be queued (4.2.1)
P26 Deal with RPM (303)
Alters no Q stores
Q4, Q5 (5.2.4)

**P27 Initial housekeeping for OUTS (5.2)
P28 Read cell count (4.1.1)
P29 Read cell count (4.1.1)
P30 Read cell count (4.1.1)
P31 Fetch character to 07 (4.2)
P32 Fetch character to 07 (4.2)
P33 Fetch character to 07 (4.2)
P34 Fetch character to 07 (4.2)
P35 Fetch character to 07 (4.2)
P36 Fetch character to 07 (4.2)
P37 Fetch character to 07 (4.2)
P38 Fetch character to 07 (4.2)
P39 Fetch character to 07 (4.2)
P40 Fetch character to 07 (4.2)
P41 Fetch character to 07 (4.2)
P42 Fetch character to 07 (4.2)
P43 Fetch character to 07 (4.2)
P44 Fetch character to 07 (4.2)
P45 Fetch character to 07 (4.2)
P46 Fetch character to 07 (4.2)
P47 Fetch character to 07 (4.2)
P48 Fetch character to 07 (4.2) (also used in Phase 0)
P49 Fetch character to 07 (4.2)
P50 Fetch character to 07 (4.2)
P51 Fetch character to 07 (4.2)
P52 Fetch character to 07 (4.2)
P53 Fetch character to 07 (4.2)
P54 Fetch character to 07 (4.2)
P55 Fetch character to 07 (4.2)
P56 Obtain program letter, given priority number
P57 Obtain priority number given program letter in C7 (402)
P58 Find tape with given identifier (4.1.6)
P59 Update Flexowriter queue (4.1.3)
P60 Deal with RPM (303)
Alters no Q stores
Q4, Q5 (5.2.4)

**P61 Initial housekeeping for OUTS (5.2)
P62 Controller of W mag tape bricks (4.1.1)
P63 W mag tape bricks (4.1.1)
P64 Hold up subprogram before transfer
P65 Set up Q7 with program parameters
Alters Q6, Q7

P66 Read cell count (4.1.1)
P67 Read cell count (4.1.1)
P68 Read cell count (4.1.1)
P69 Read cell count (4.1.1)
P70 Read cell count (4.1.1)
P71 Read cell count (4.1.1)
P72 Read cell count (4.1.1)
P73 Read cell count (4.1.1)
P74 Read cell count (4.1.1)
P75 Read cell count (4.1.1)
P76 Read cell count (4.1.1)
P77 Read cell count (4.1.1)
P78 Read cell count (4.1.1)
P79 Read cell count (4.1.1)
P80 Read cell count (4.1.1)
P81 Read cell count (4.1.1)
P82 Read cell count (4.1.1)
P83 Read cell count (4.1.1)
P84 Read cell count (4.1.1)
P85 Read cell count (4.1.1)
P86 Read cell count (4.1.1)
P87 Read cell count (4.1.1)
P88 Read cell count (4.1.1)
P89 Read cell count (4.1.1)
P90 Read cell count (4.1.1)
P91 Read cell count (4.1.1)
P92 Read cell count (4.1.1)
P93 Read cell count (4.1.1)
P94 Read cell count (4.1.1)
P95 Read cell count (4.1.1)
P96 Read cell count (4.1.1)
P97 Read cell count (4.1.1)
P98 Read cell count (4.1.1)
P99 Read cell count (4.1.1)
P100 Read cell count (4.1.1)
P101 Read cell count (4.1.1)
P102 Read cell count (4.1.1)
P103 Read cell count (4.1.1)
P104 Read cell count (4.1.1)
P105 Read cell count (4.1.1)
P106 Read cell count (4.1.1)
P107 Read cell count (4.1.1)
P108 Read cell count (4.1.1)
P109 Read cell count (4.1.1)
P110 Read cell count (4.1.1)
P111 Read cell count (4.1.1)
P112 Read cell count (4.1.1)
P113 Read cell count (4.1.1)
P114 Read cell count (4.1.1)
P115 Read cell count (4.1.1)
P116 Read cell count (4.1.1)
P117 Read cell count (4.1.1)
P118 Read cell count (4.1.1)
P119 Read cell count (4.1.1)
P120 Read cell count (4.1.1)
P121 Read cell count (4.1.1)
P122 Read cell count (4.1.1)
P123 Read cell count (4.1.1)
P124 Read cell count (4.1.1)
P125 Read cell count (4.1.1)
P126 Read cell count (4.1.1)
P127 Read cell count (4.1.1)
P128 Read cell count (4.1.1)
P129 Read cell count (4.1.1)
P130 Read cell count (4.1.1)
P131 Read cell count (4.1.1)
P132 Read cell count (4.1.1)
P133 Read cell count (4.1.1)
P134 Read cell count (4.1.1)